

WHITE PAPER

Server Innovations: Examining DC Power as an Alternative for Increasing Data Center Efficiency and Reliability

Sponsored by: Rackable Systems Inc.

John Humphreys

Jessica Yang

August 2004

IDC OPINION

In today's cost-conscious yet mission-critical IT world, IT directors and managers are increasingly pressed to examine their data centers and find economies of scale as well as ways to increase data center efficiency. Some recent trends have included the wholesale migration to scale-out environments as well as data center and/or server consolidation as a means to lower both capital and operational costs. These initiatives have helped users become more efficient and productive, but they address only a portion of the issues at hand.

Driven by the high costs of power and exhausted budgets, some data centers are considering DC power options to satisfy their need for reliable infrastructure. This emerging trend in the commercial world is both a result of these consolidation and scale-out trends and a way to help users increase the reliability of services running on that infrastructure. IDC believes the trend of DC power in the data center can help users to cut costs and improve the reliability of their critical IT environments.

IN THIS WHITE PAPER

This white paper examines the relationship between power consumption and overall data center costs and the impact it will have as organizations seek to build out their infrastructures while reducing overall cost of service. It also looks at the ways in which increasing power requirements are accelerating the demand for better thermal management.

This document provides an overview of Rackable Systems' DC power solutions as technology that provides increased reliability and more efficient thermal management in large-scale data center environments. It profiles the experiences to date of a customer of the new technology, Toshiba America Electronic Components Inc. (TAEC). This paper also examines some of the challenges and opportunities that face Rackable Systems in bringing this new solution into the marketplace.

SITUATION OVERVIEW

While IT organizations are under increasing pressure to deliver higher levels of service to end users, they also are faced with the need to reduce costs. Companies of all sizes and types are struggling as they try to balance increasing infrastructure demands and changing business practices with reduced IT budgets.

Maintaining reliability and availability is at the top of every IT professional's list of priorities. Nothing is more fundamental to keeping a data center environment functioning than the power and cooling supplied to the environment. Power resources must be consistent and reliable to meet the needs of even small data centers. In larger, more complex data centers supporting mission-critical applications, the need for efficient power and reliable servers is further amplified.


IDC has found that only 56% of total IT costs are related to hardware and software; the rest are administration and overhead costs. Power and cooling constitute a significant part of the operational spending in data centers. As a result, IT organizations can better control system deployment costs through reducing power consumption expenses as well as related management and maintenance costs.

IT managers' increasing adoption of faster generations of microprocessors is compounding the power concerns. As businesses grow, organizations will continue to build out their IT infrastructures to support more sophisticated applications and external integrations. However, processors generate much of the heat in the data center, and the deployment of more systems with faster processing speeds will naturally increase both power consumption and heat output. The ability to sustain the next generation of microprocessors is largely contingent on the ability to effectively manage heat, both input and output, and use power as efficiently as possible.

IDC forecasts that the number of servers shipped will increase at a compound annual growth rate of 17% over the next five years (see Figure 1). Many of these investments will focus on the scale-out computing environment, or the deployment of multiple, smaller servers, thus increasing the number of systems present in the data center.

FIGURE 1

Historical and Forecast Server Shipments, 2000–2008




Source: IDC, 2004

In addition, there is a trend toward consolidation in data centers with clustered and rack-mounted systems, which leads to the deployment of more hardware in a fixed amount of space to increase the server density. Clustered systems are commonly deployed in the industry, driven by their support for highly available applications and data, the ability to flexibly support growing workloads by adding server nodes, and the ability to support multiple servers through unified system management.

Over the past three years, rack-optimized systems have made up an increasing share of the servers deployed in data centers. IDC expects this trend to continue as rack-mount and blade servers become more pervasive in data centers. IDC forecasts that average server density will go from 3U in 2003 to 1U by 2008, effectively tripling server density over a five-year period (see Figure 2).

FIGURE 2

Historical and Forecast Adoption of Pedestal, Rack-Optimized, and Blade Servers, 2000–2008



Source: IDC, 2004

Increasing server density brings additional power concerns. As the number of servers increases, more power must be supplied to keep the servers operating, which results in higher power consumption and costs. Increased density also means greater heat production from the processors and power supplies, the two primary generators of heat in the data center environment. This additional heat generated, in turn, requires more power to cool the data centers, which drives up the associated electricity cost.

Additional heat can also decrease the reliability of the servers. IT staffs spend more time responding to downtime and deploying replacement systems, which leads directly to increased system administration expenses. Higher operating temperatures cause systems to prematurely reach the end of their lives, resulting in greater hardware costs. Other costs are associated with productivity loss as users experience disruption in service due to system failure.

All of these factors accentuate the importance of efficient, reliable power resources that are critical to maintaining cost-effective, flexible, and dependable data centers. IDC research indicates that a growing number of IT organizations are moving toward a model of "service-centric" computing, where IT services, storage, and network resources are made available to business units or end users in a manner that is fully accountable, metered, and always available. In such a model, IT managers will focus on the services delivered to ensure that there is no single point of failure and to provide high availability. As a result, IT organizations will increasingly consider the impact of power consumption and costs and look for ways to gain efficiency and reliability while reducing overall costs as they plan future investments.

RACKABLE SYSTEMS PRODUCTS

Rackable Systems has designed an innovative solution to address power management problems facing data centers today. Leveraging DC power distribution, Rackable Systems refines and applies new techniques to existing technologies, bringing DC power solutions into its back-to-back rack-mount servers.

A DC power supply, which is offered as an option on Rackable Systems servers, replaces the standard AC power supply inside the servers. This DC power distribution strategy is more efficient than traditional AC designs, particularly for the high wattages required by the latest processor technologies. For instance, 350W AC power supplies are often between 65% and 75% efficient; therefore, a server that needs 230 watts would draw about 320 watts of external AC power. In comparison, Rackable Systems' solution, with its DC cabinet powered by 48V DC data center power, is 93% efficient at the server level, which means that the same server would draw only 250 watts of external power. On this factor alone, the DC power technology reduces server power consumption by nearly a third.

Because the solution is compatible with the common AC-based data centers, a telco environment built upon DC power is not required for customers to deploy a Rackable Systems DC power solution. Rectifiers that convert incoming AC power to DC power are placed either at the top of the cabinet or externally near an HVAC unit. Data centers providing DC power directly to the cabinet will see the largest power savings, up to 35% over AC-based systems, depending on the comparison points. Lower power requirements at the server level can mean substantial reduction in a data center's monthly electricity bill, especially for large-scale deployments. Furthermore, Rackable Systems' server efficiency allows for up to 30% higher server density in data centers with a fixed amount of power and cooling available per square foot.

The ability to sustain faster processors hinges upon being able to properly cool servers and to keep data centers at a safe and constant temperature range. Watts of power used by a server are dissipated into the air as heat, and the two components that generate the most heat are processors and power supplies. Rackable Systems' DC power solutions consume less power and are able to shift a large heat load out of the servers to the rectifiers. Such a design generates between 20% and 40% less heat at the system level, compared with the design of a traditional AC-powered rack in a typical AC power data center. Lower heat output reduces the cost of air conditioning needed to cool data centers — another way to save on monthly electricity expenses.

As each DC power system uses fewer watts and runs cooler, administrators are less likely to experience frequent heat-related failures and IT staff will spend less time tending to heat-worn components. Rackable Systems' DC power supply provides a mean time between failure (MTBF) rate of more than 7 million hours, whereas standard AC power supplies have an MTBF of only 100,000 hours. This 70-times improvement in reliability at the power supply level translates to increased overall server reliability. Server uptime is further enhanced as individual servers receive N+M redundant power via cabinet- or data center-level rectifiers delivered through DC distribution bars. Redundancy at the rectifier level means system power can be designed to survive without impact in the event of an AC circuit failure.

Rackable Systems' DC power solutions enable data centers to achieve the reliability, flexibility, and manageability that are critical today. Improving overall server reliability ensures high accessibility of the most mission-critical applications. Reducing power and administration costs lowers the total cost of ownership. Rackable Systems' customers can leverage the next generation of processors and benefit from the higher performance knowing that their investments are protected.

CASE STUDY

An early customer of Rackable Systems' DC power solution is Toshiba America Electronic Components Inc. (TAEC), a semiconductor design company with several design centers in the United States. The company began an upgrade to its 32-bit Linux server environment in 2003, with the goal of adding compute resources to handle increasing design analysis workloads. With a fixed floor space in its data centers, TAEC had to deploy clusters that would maximize density with dual processors in a 1U rack-mount configuration. Rackable Systems' patented half-depth servers, which mount back-to-back in a standard rack, offered the density and added performance that TAEC required.

Deploying Rackable Systems' back-to-back rack-mount servers satisfied TAEC's critical requirements for more compute power at the highest possible density. But with the increased processing speeds and density comes the obvious challenge of managing greater power consumption and thermal output. TAEC required a complete solution that would leverage its existing power and HVAC infrastructure while providing the added compute power and density without dramatically increasing its electrical costs.

Rackable Systems' DC power technology satisfied that additional requirement. With DC power-based servers, TAEC was able to deploy high-density clusters with maximum processing power without concern for increased electricity bills or higher server failure rates. Because Rackable Systems' DC power solutions are more efficient, TAEC was able to add more compute resources and achieve better price performance ratios without increasing IT costs. TAEC's data centers are managed by the same number of IT resources in spite of a dramatic increase in the number of systems — since less time is spent on maintenance and system failures.

In addition to lower initial hardware acquisition cost and ongoing power savings, TAEC credited Rackable Systems' superior service and comprehensive warranty as differentiating factors in its decision. The maintenance program and attentive sales and support staff have made owning Rackable Systems' products a worry-free investment for TAEC. The company was able to quickly deploy its next-generation EDA applications to two design centers in San Jose, California, and Marlborough, Massachusetts, because it was able to customize systems to its exact configuration needs. Because Rackable Systems' servers come fully configured, installation was as simple as plugging in the rack, and only minimal in-house integration was required.

Like many industries, the semiconductor industry is keenly focused on time to market. This requires TAEC to invest heavily in its IT infrastructure to keep pace with the latest server and processor technologies available. Rackable Systems' patented highly dense servers — with the added benefit of DC power technology — have made it possible for TAEC to double the number of servers in its data centers while keeping IT costs in check to meet its critical business needs.

CHALLENGES/OPPORTUNITIES

Rackable Systems has introduced its DC power solution at a time when reducing the elemental costs associated with the data center — such as power consumption — is an important priority for IT professionals. As IT budgets expand and organizations consider new investments, Rackable Systems should find organizations that are receptive to discussions about increasing the efficiency of their power resources. In addition, market demand for 64-bit processors will accelerate to support more compute-intensive applications, opening the door for discussions about power consumption.

Rackable Systems faces several challenges that are centered on raising market awareness and establishing credibility. It must invest in creating broad market awareness of both the problems associated with power consumption in data centers and the option of using DC power as a solution. Most IT professionals already know that power is a significant concern, but Rackable Systems could benefit from raising their understanding of the compounding impacts of power, cooling, and increased density in relation to next-generation processors as they plan their infrastructure buildouts.

The lack of awareness among potential customers about the applicability of the DC power solution is another challenge. DC power is a well-understood technology, especially in the telco industry, but Rackable Systems is applying it in new ways. The company will need to educate IT professionals on the benefits and impact of deploying this new solution. Rackable Systems has made progress on this front already, winning the 2004 Linux World New York Conference & Expo Best Cluster Solution Award. The company must continue to gain such visibility and to highlight the experiences of its early adopters.

If the downside to having a new type of solution is the need to establish credibility, the upside is being the only provider in the marketplace. If Rackable Systems is able to build demand for its solution and establish its credibility, it will benefit from being the first and only vendor with this unique technology.

CONCLUSION

As the business climate improves and IT spending resumes, companies of all types and sizes will look to build their infrastructure environments with added flexibility and lower management and operating costs. Increasing deployments of multiple smaller systems and continuing server and storage consolidation contribute to higher CPU density in data centers today. Combined with escalating power draws on next-generation processors, the need for a power efficient infrastructure solution has never been more pressing. Systems availability is contingent upon the ability to remove heat from servers and increase the reliability of power supplies.

DC power has been the power standard in telco markets for quite some time, but only recently is DC input becoming a more popular option for new data centers. Now Rackable Systems has applied this mature technology in an innovative way to address the accelerating demand for power consumption. DC power supplies consume fewer watts of power and generate less heat compared with traditional AC power supplies. This increased power efficiency not only lowers electricity bills but also allows for high-density server environments with the latest processor technology.

DC power solutions help to extend the MTBF for server hardware and reduce the time IT administrators spend attending to failed hardware, which directly lowers system management expenses and hardware replacement costs for heat-worn components. With improved reliability, companies can avoid costly server downtime and any associated productivity loss.

IT organizations are moving toward a model of "service-centric" computing and coming under increasing pressure to provide higher levels of service and availability to their end users. As IT managers plan future investments, they will pay more attention to the cost of power consumption and the need for a more efficient and reliable power solution. Rackable Systems' DC power solution is a proven technology and will enable users to achieve the goal of a cost-effective, flexible, and reliable data center environment.

Copyright Notice

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2004 IDC. Reproduction without written permission is completely forbidden.